A FPGA-Based Fast Transient Search Platform for the Australian Square Kilometre Array Pathfinder

Robert Navarro[(1)], Nathan Clarke [(2)], Larry D'Addario*[(1)], Joseph Trinh[(1)] and Tsan-Huei Cheng[(1)]

(1) Jet Propulsion Laboratory, California Institute of Technology

(2) International Centre for Radio Astronomy Research

Background

- Parameters of the ASKAP telescope
  - N = 36 antennas, each a 15m diameter dish
  - K = 15 to 36 dual-pol beams per antenna, formed by phased array feeds
  - f = 700 to 1800 MHz tuning range; 304 MHz processed bandwidth
  - B = 1 MHz channel bandwidth (coarse filter bank, 304 channels)
  - Configuration: 27 antennas in 1 km radius core, max baseline 6 km.
- Commensal Real-time ASKAP Fast Transients (CRAFT) survey
  - Commensal survey for fast (<<5 s) transient signals
  - One of 10 survey science projects selected for ASKAP by international review of proposals. CRAFT team has 42 members from 14 institutions (9 outside Australia).
  - Science: giant pulses, rotating radio transients (RRATs), magnetars, gamma-ray bursts, annihilating black holes, gravitational waves, SETI.
  - Search in dispersion measure over DM = 10 to 3000 pc/cm^3 by computing de-dispersed time series for many DMs in parallel
  - Make use of coarse power spectra (304 channels) available for each beam of each antenna every ~1 ms.

ASKAP Processing – Simplified View

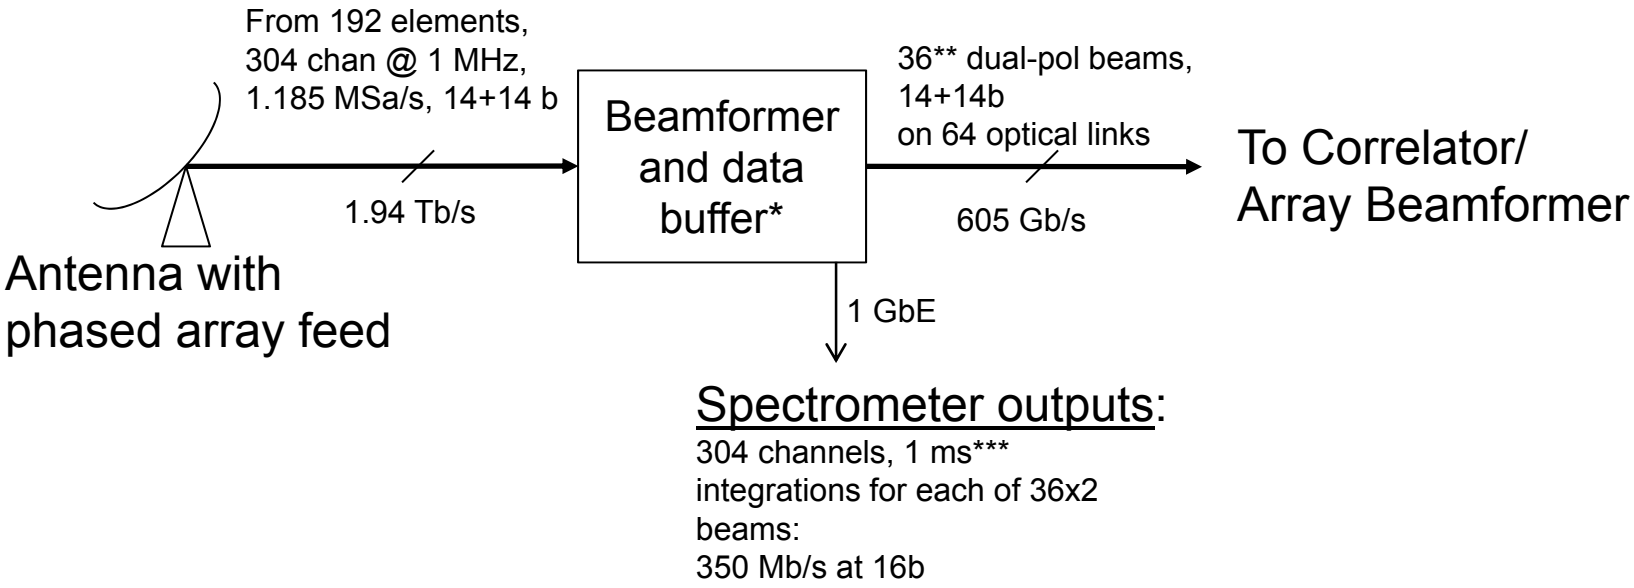

Notes:

* **Beamformer contains circular buffer for capturing input samples.**

** Maximum of 36 beams at high frequency end of tuning range, fewer at low frequency end.

*** Shorter integration times may be possible.

CRAFT Transient Search Back End

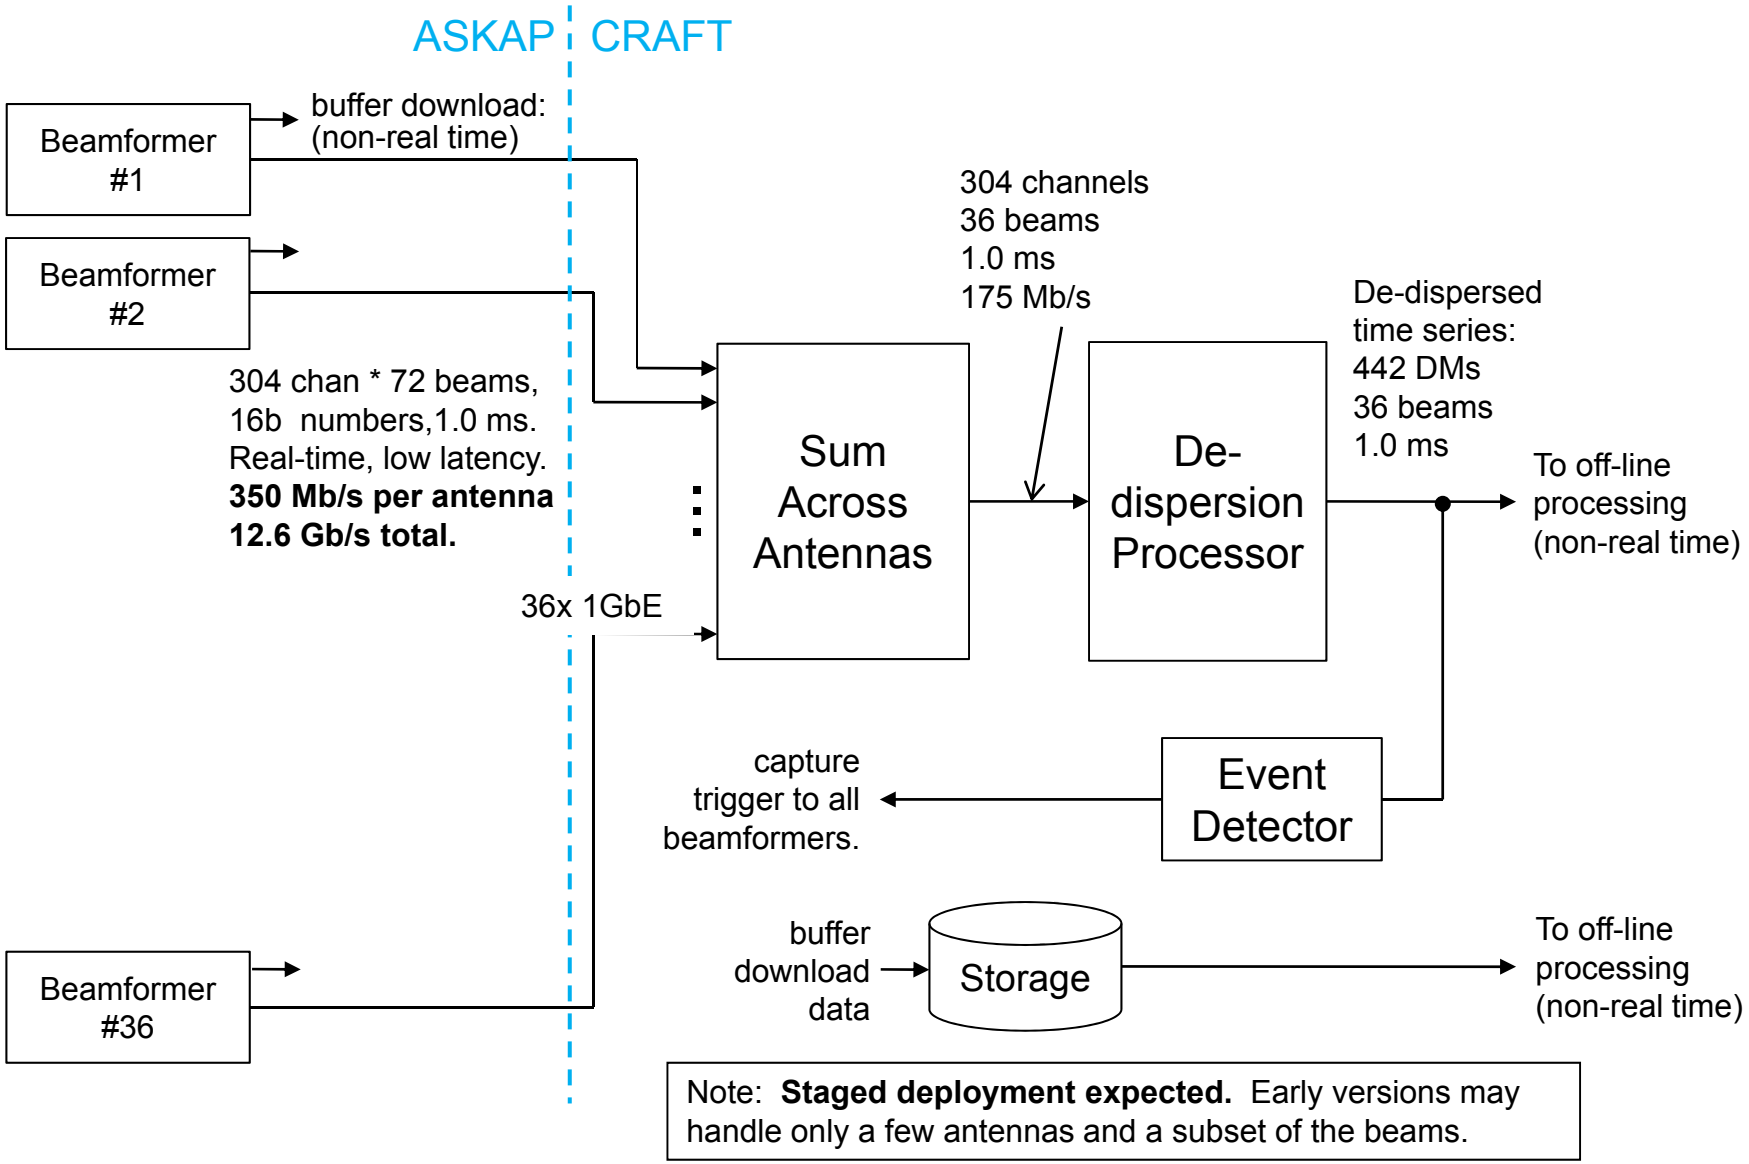

Implementing the CRAFT Back End

- Two main processes in CRAFT back end:
  - Summing of spectrometer outputs across antennas
  - Parallel de-dispersion for multiple DMs and multiple beams.
  - Requires real-time processing with low latency.
- Implementation approaches
  - Our team at JPL and Curtin University is investigating an FPGA-based implementation.
  - Another team lead by Richard Dodson at University of Western Australia is investigating GPU- based implementations.
  - This talk concentrates on the FPGA-based implementation, but many concepts apply to both.

Coherent vs. Incoherent combining

- With *N* antennas, signals can be combined by
  - Cross-correlation (coherent)
    - Allows imaging, but practical time resolution is inadequate for transients (minimum integrating time for the ASKAP correlator is 5 s)
  - Summation in a beam former (coherent)
    - Highest sensitivity
  - Post-detection summation (incoherent)
    - Largest field of view
- Which gives highest survey speed? For identical antennas:
  - incoherent: N where *N* is the number of antennas.
  - coherent: $b\eta N^2$ where b is the number of array beams and η is the antenna packing density.
  - ASKAP: N=36, b=1, η=1.6x10^{-5} $\Rightarrow$ 36 incoherent, .021 coherent
  - SKA: N=660 within 1 km, η=2.6x10^{-4} $\Rightarrow$ coherent is better for b>6.
  - See SKA Memo 123.
- Therefore, incoherent combining was chosen for CRAFT

Dispersed Pulse Received by a Filter Bank Spectrometer & SNR Loss When DM is Wrong

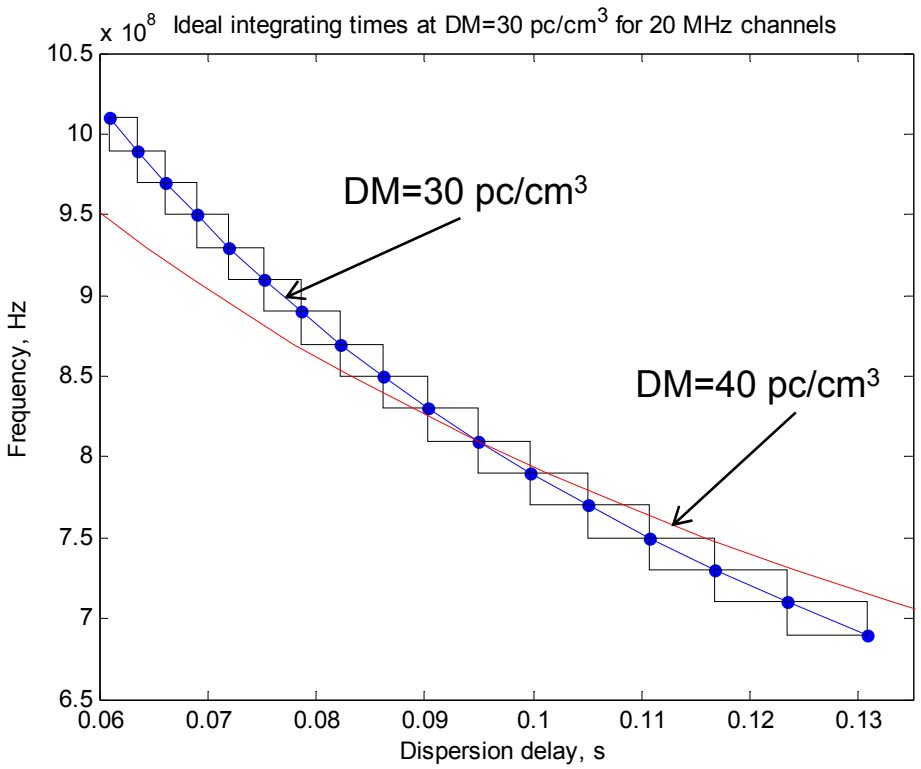

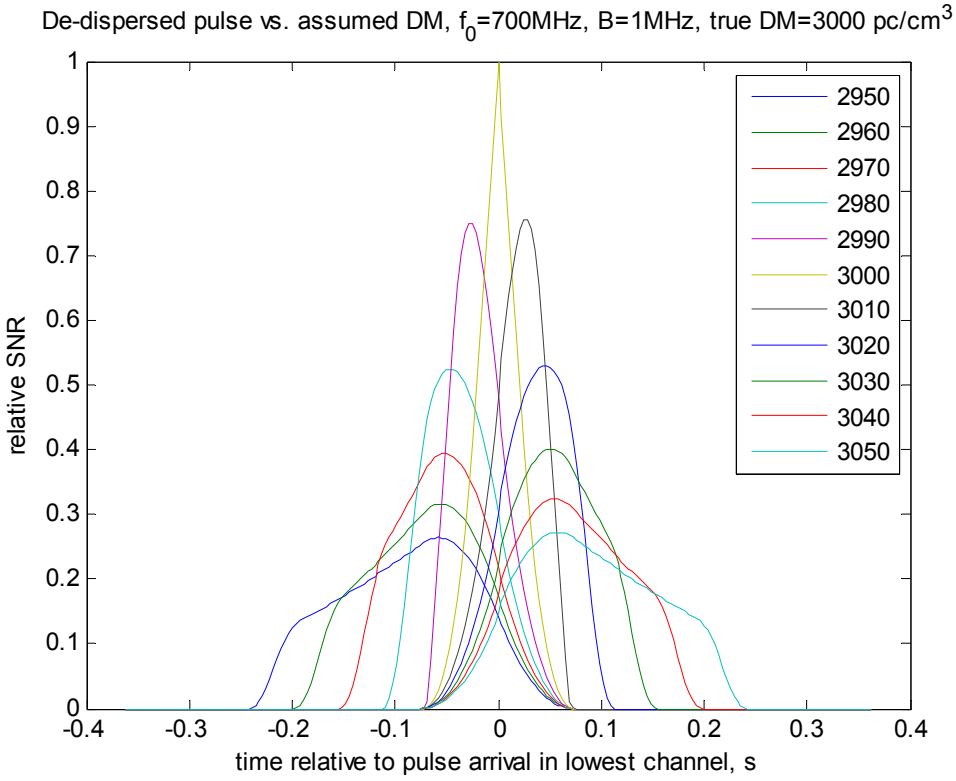

How many DMs must be searched?

- Incoherent de-dispersion used: Delay and sum across channels.
- For ASKAP's parameters (304 channels, 700 to 1004 MHz), an error of 0.67% in DM reduces peak SNR by 1/2 (−3dB); see previous slide and SKA Memo 124.
- Therefore, searching DM in steps of r = 1.3% ensures that a pulse whose DM falls between search values will be detected at no less than 1/2 of its true amplitude.
- This means that search values increase exponentially with step number.
- To cover the range DM = c_1 to c_2 requires

  $$\log(c_2/c_1)/\log r$$

  steps.
- DM = 10 to 3000 pc/cm^3 then requires 442 steps.
- However:
  - At high DM, pulse broadening from scattering can allow coarser steps
  - DM range 10 to 3000 pc/cm^3 is a straw man. Some people seem to think we need at least that much, others think it's too much.

Processing Requirements

- Assume these parameters
  - DM range= 10 to 3000 pc/cm^3
  - Beamformer spectrometer outputs every 1.0 ms.
  - 304 channels of 1 MHz BW (in freq range of 1.8 to 0.7 GHz)
- Incoherent de-dispersion requires only additions.
- Rate of additions for de-dispersion
  - Per Beam: 1.0x10^6 adds to compute 1 de-dispersed sample for each of the 442 DMs (at low end of freq range)
  - Do this every 1.0 ms ⇒ 1.0x10^9 adds/second.
  - Per 36 Beams: 36.0x10^9 adds/s (at low end of freq range)
- Memory size
  - DM=3000 pc/cm^3 requires ~13s of samples in buffer.
  - Per Beam: 7.9 Mbytes (for 16b values). 285 MB for 36 beams.
- Memory bandwidth
  - Writes to buffer: 0.6 Mbytes/s per beam (incoming data)
  - If each add requires reading a value from the buffer, then
    Reads from buffer: 2.0 Gbytes/s per beam, 72 GB/s for 36 beams.

Implications of Processing Requirements

- De-dispersion arithmetic is not much of a challenge for FPGA
  - Need 181 adders running at 200 MHz in FPGA to handle de-dispersion processing for 36 beams
  - One Xilinx XC6V240T has 768 DSP slices containing add-accumulate capability.
- Buffer memory size exceeds on-chip FPGA capacity
  - need off-FPGA DRAM.
- Memory bandwidth issues
  - For all 36 beams, memory bandwidth of 72 Gbytes/s might be needed.
  - A 64 bit DDR3 DIMM has a nominal bandwidth of ~5 Gbytes/sec.
  - Achieving high DRAM bandwidth is strongly dependent on data organization; practical bandwidth may be 2x to 8x less than nominal.
  - Therefore, de-dispersion is memory bandwidth limited.
  - We have studied several approaches to reducing the memory bandwidth by avoiding reading the same value multiple times (next slide).

Improving Memory Bandwidth for De-dispersion

- Approaches:
  1. Of the 1 million samples in the buffer that are needed to compute one de-dispersed sample for each of 442 DMs, many are used for more than one DM and could be read only once.
  2. If we don't compute new de-dispersed samples every ms but instead compute a batch of, say, 32 of them all at once, then most of the buffer stays the same and need be read only once.

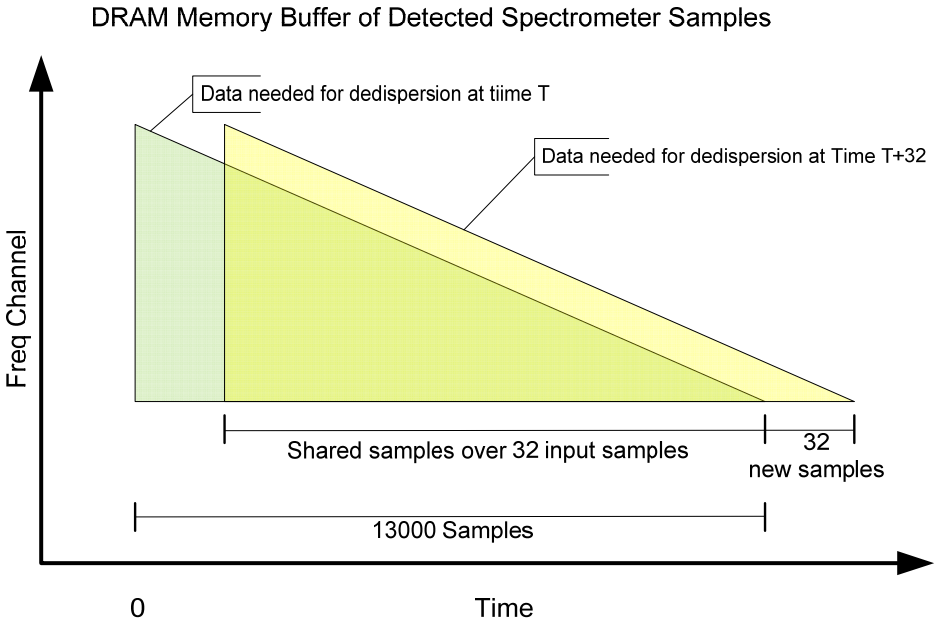

- Approach 2 above produces the biggest bandwidth improvement.
- For one frequency channel at a time, all 13032 time samples are read into FPGA memory and the appropriate ones are added to each of 32*442 output accumulators.
- This requires DRAM bandwidth of 5.3 Gbytes/s for 36 beams (versus 72 Gbytes/s for new outputs every ms).
- Down side: adds latency to processing. De-dispersed output is delayed up to 32 ms.
  - Minimization of latency is important for triggering the capture buffer.
  - Capture buffer length of >1 s is expected, so 32 ms seems acceptable.
- Current FPGA implementation uses 4 XC6V240T FPGAs, each with 32 bit DRAM.
  - This algorithm fits comfortably. The present design uses only 16-sample output blocks.

CRAFT FPGA Hardware

- Implemented on COTS PCIe card with plug in FPGA modules
  - Pico Computing Ex-500 backplane and M-501 modules.
  - One FPGA module for cross-antenna summing, four for de-dispersion.
  - Each module has one Virtex-6 LX240T-2 and 512 MB DDR3 memory.
  - Backplane has x16 Gen2 PCIE to host, 8GB/s bandwidth
- Data from beamformers received over dual 10GbE transceivers.
  - Myricom Dual SFP+ 8x PCIE Gen2 network interface card

EX-500 backplane with two M-501 FPGA modules installed.

- Also Considered:
  - Casper Roach Board
  - ASKAP Redback Board.

Implementation Details/Status for Pico FPGA

- Logical organization of processing shown below.
- Detailed implementation of de-dispersion blocks shown in gray is in progress.
- One FPGA module for antenna summations, 4 FPGA modules for de-dispersing all 36 beams. System fits on one PCIe board.

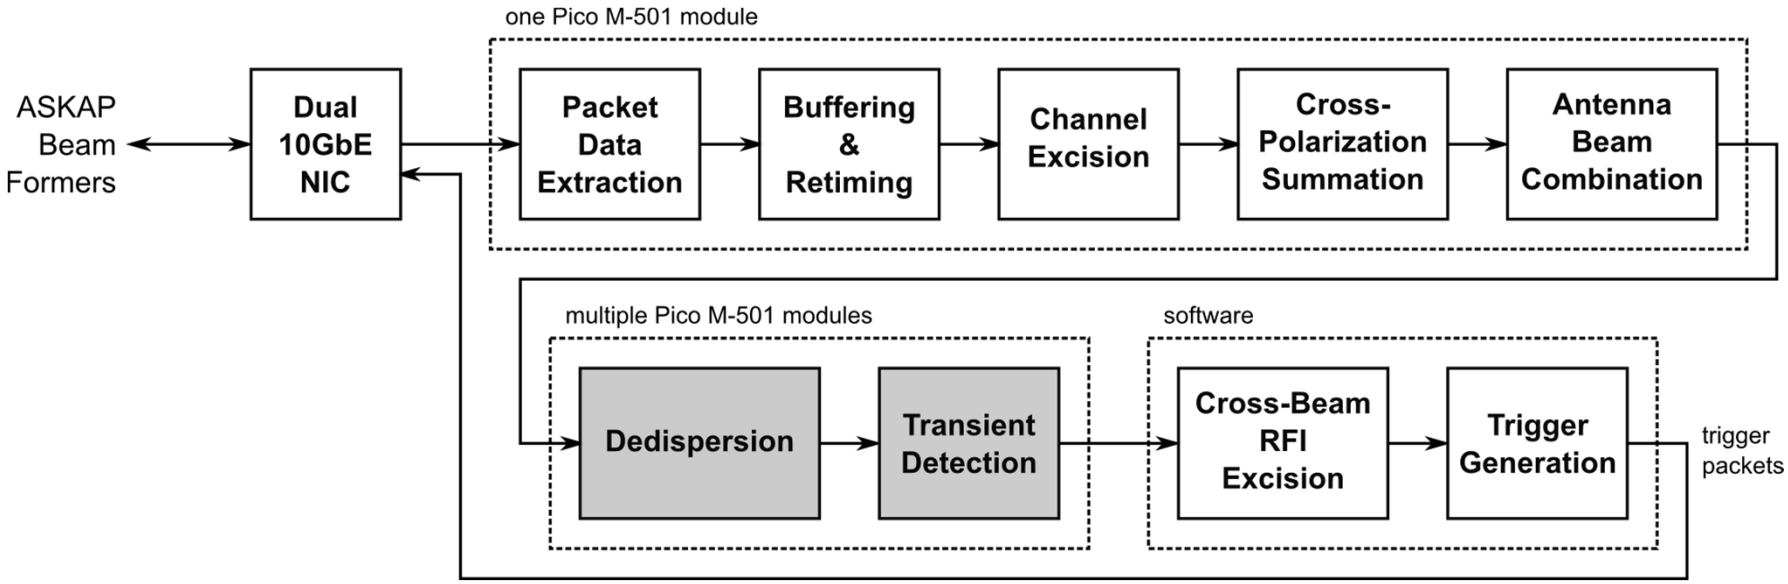

Conclusion

- A fast transient detection backend is desired by the CRAFT survey for the ASKAP telescope.
- Analysis of the science goals and design requirements suggest that incoherent summation across antennas and incoherent de-dispersion searches are best with ASKAP – may not be so for full SKA.
- FPGAs have more than enough processing power to perform the de-dispersion summations for DM=10 to 3000 pc/cm^3 for all 36 beams.
- Memory bandwidth, rather than arithmetic, limits how much can be done in each FPGA.
- Trading latency for bandwidth enables a workable design.
- PCIe based COTS board chosen for FPGA implementation.
- Detailed Verilog coding of de-dispersion algorithm in progress.